

LEARNING AND EVOLUTIONARY GAME THEORY

Luis R. Izquierdo
University of Burgos
Burgos
Spain
luis@izquierdo.name

Segismundo S. Izquierdo
University of Valladolid
Valladolid
Spain
segis@eis.uva.es

Fernando Vega-Redondo
European University Institute
Florence
Italy
Fernando.Vega@eui.eu

Definition

Game: a formal abstraction of a social interaction where: (a) there are two or more decision makers, called players, (b) each player has a choice of two or more ways of acting, called actions or (pure) strategies, and (c) the outcome of the interaction depends on the strategy choices of all the players.

Game theory: the formal theory of interdependent decision-making.

Classical Game Theory: branch of game theory devoted to the formal analysis of how rational players should behave in order to attain the maximum possible payoff.

Evolutionary Game Theory: branch of game theory that studies the evolution of large populations of individuals who repeatedly play a game and are exposed to evolutionary pressures (i.e. selection and replication subject to mutation).

Learning Game Theory: branch of game theory that studies the dynamics of a group of individuals who repeatedly play a game, and who adjust their behavior (strategies) over time as a result of their experience (through e.g. reinforcement, imitation, or belief updating).

Theoretical Background

This entry gives a comparative overview of Learning and Evolutionary Game Theory within the broader context of non-cooperative game theory. We make a clear distinction between game theory used *as a framework* (which makes no assumptions about individuals' behavior or beliefs), and the three main branches of non-cooperative game theory as we know them nowadays, namely classical game theory, evolutionary game theory and learning game theory.

Game theory as a *framework* is a methodology used to build models of real-world social interactions. The result of such a process of abstraction is a formal model (called *game*) that typically comprises the set of individuals who interact (called *players*), the different choices available to each of the individuals (called actions or *pure strategies*), and a *payoff* function that assigns a value to each individual for each possible combination of choices made by every individual (Fig. 1). In most cases, payoffs represent the preferences of each individual over each possible outcome of the social interaction. The notable exception is evolutionary game theory, where payoffs often represent Darwinian fitness.

		Player 2	
		Player 2 chooses LEFT	Player 2 chooses RIGHT
Player 1	Player 1 chooses UP	3, 3	0, 4
	Player 1 chooses DOWN	4, 0	1, 1

Fig. 1. Payoff matrix of a 2-player 2-strategy game. For each possible combination of pure strategies there is a corresponding pair of numbers (x, y) in the matrix whose first element x represents the payoff for player 1, and whose second element y represents the payoff for player 2.

Game theory is particularly useful to model social interactions where individuals' decisions are interdependent, i.e. situations where the outcome of the interaction for any individual player generally depends not only on her own choices, but also on the choices made by every other individual. Thus, several scholars have pointed out that game theory could well be defined as “the theory of interdependent decision-making”.

Game theory *used as a framework* provides a formal description of the social setting where the players are embedded. Importantly, it does not account for the players' behavior, neither in a normative nor in a positive sense. It is just not the realm of game theory *as a framework* to do so. It is only when different assumptions about the precise meaning of payoffs and about how players behave –or should behave– are included in the framework, that game theory gives rise to its different branches. Here we outline the main features of the three most developed branches of non-cooperative game theory at this time:

Classical game theory (CGT): Classical game theory was chronologically the first branch to be developed (Von Neumann and Morgenstern 1944), the one where most of the work has been focused historically, and the one with the largest representation in most game theory textbooks and academic courses.

In CGT payoffs reflect preferences, i.e. the payoffs for each player effectively define a preference ordering over the set of possible outcomes. Naturally, the specific properties of this ordering constrain the type of analysis that one can meaningfully undertake. Thus, the most basic assumption about payoffs is to assume that they merely represent a total ordering of preferences, e.g. “Worst, Average, Best” (so arithmetic operations on payoffs would not be meaningful). However, most often payoffs in CGT are interpreted as von Neumann-Morgenstern utilities and, when this is the case, payoffs embody players' attitudes to risk, and thus one can use expected utility theory to evaluate probability distributions over possible outcomes of the game. This allows for the analysis of *mixed strategies*, which are strategies that assign a certain probability to each possible pure strategy. Importantly, note that if no further assumption is made, comparisons of payoffs across players are completely meaningless. (This contrasts with the interpretation of payoffs made in evolutionary game theory, where payoffs represent reproduction or survival rates and it is the relative differences in payoffs among players what actually drives the dynamics of the process.)

In CGT players are assumed to be rational, meaning that they act as if they have consistent preferences and unlimited computational capacity to achieve their well-defined objectives. The aim of the theory is to study how these instrumentally rational players would behave in order to obtain the maximum possible payoff in the formal game. The main problem in CGT is that, in general, assuming rational behavior for any one player

rules out very few actions, and consequently very few outcomes, in the absence of strong assumptions about other players' behavior. Hence, in order to derive specific predictions about how rational players would behave, it is often necessary to make very stringent assumptions about everyone's beliefs and their reciprocal consistency. With these assumptions in place, the outcome of the game is a Nash equilibrium, which is a set of strategies, one for each player, such that no player, knowing the other players' strategies in that set, could improve her expected payoff by unilaterally changing her own strategy.

Given the strength of the assumptions usually made in CGT, it is not surprising that when game theoretical solutions have been empirically tested, disparate anomalies have been found. To make matters worse, even when the most stringent assumptions are in place, it is often the case that several possible outcomes are possible, and it is not clear which –if any– may be achieved, or the process through which this selection would happen. Thus, in general, the direct applicability of CGT is limited. A related limitation of CGT is that it is an inherently static theory: it is mainly focused on the study of end-states and possible equilibria, paying hardly any attention to how such equilibria might be reached.

Evolutionary Game Theory (EGT): Some time after the emergence of classical game theory, biologists realized the potential of game theory as a framework to formally study adaptation and coevolution of biological populations, particularly in contexts where the fitness of a phenotype depends on the composition of the population (Hamilton 1967). The main assumption underlying evolutionary thinking is that the entities which are *more successful* at a particular time will have the best chance of being present in the future. In biological and economic contexts, this assumption often derives from competition among entities for scarce resources or market shares. In other social contexts, evolution is often understood as *cultural* evolution, and it refers to dynamic changes in behavior or ideas over time.

In general, a model is termed evolutionary if its laws of motion reflect the workings of three mechanisms: *selection*, *replication*, and *mutation*, appropriately interpreted for the context in hand. The mechanism of *selection* is a discriminating force that favors some specific entities rather than others. Within the context of game theory, this selection is based on payoffs, so players that have obtained higher payoffs are selected preferentially over those with relatively lower payoffs. The *replication* mechanism ensures that the properties of the entities in the system (or the entities themselves) are preserved, replicated or inherited from one generation to the next at least to some extent. Within the context of evolutionary game theory, the *replication* mechanism ensures that the strategies of selected players are adequately inherited, or transmitted, across consecutive generations. *Selection* and *replication* are two mechanisms that work very closely together, since *being selected* means *being selected to be preferentially replicated*.

In general, the workings of selection and replication tend to reduce the diversity of the system. The generation of new diversity is the job of the *mutation* mechanism, which is a process that works alongside (and in opposition to) the homogenizing mechanisms of selection and replication to preserve the heterogeneous nature of the system, i.e. the everlasting presence of different strategies. This mutation process by which new entities or new patterns of behavior appear is often called experimentation or innovation in socio-economic contexts.

EGT is devoted to the study of the evolution of strategies in a population context. In biological systems, players are typically assumed to be pre-programmed to play one given strategy, so studying the evolution of a population of strategies becomes formally equivalent to studying the demographic evolution of a population

of players. By contrast, in socio-economic models, players are usually assumed capable of adapting their behavior within their lifetime, switching their strategy in response to evolutionary (or competitive) pressure. However, the distinction between players and strategies is irrelevant for the formal analysis of the system in either case, since it is strategies that are actually subjected to evolutionary pressures. Thus, without loss of generality and for the sake of clarity, one can adopt the biological stand and assume that players may die and each individual player uses the same particular fixed strategy all throughout his finite life.

Thus, to sum up, EGT is devoted to the study of a population of agents who repeatedly interact to play a game. Strategies are subjected to selection pressures in the sense that the relative frequency of strategies which obtain higher payoffs in the population will increase at the expense of those which obtain relatively lower payoffs. The aim is to identify which strategies (i.e. type of players or behavioral phenotypes) are most likely to thrive in this “evolving ecosystem of strategies” and which will be wiped out by selective forces. In this sense, note that EGT is an inherently dynamic theory, even if some of its equilibrium concepts are formulated statically (e.g. the concept of evolutionary stable strategy).

In EGT, therefore, payoffs are not interpreted as preferences, but as a value that measures the success of a strategy in relation to the others; this value is often called fitness, and in biological contexts it usually corresponds to Darwinian fitness (i.e. the expected reproductive contribution to future generations). Thus, in stark contrast with classical game theory, payoffs obtained by different players in EGT will be compared and used to determine the relative frequency of different types of players (i.e. strategies) in succeeding generations. These interpersonal comparisons are inherent to the notion of biological evolution by natural selection, and pose no problems if payoffs reflect Darwinian fitness. However, if evolution is interpreted in cultural terms, presuming the ability to conduct interpersonal comparisons of payoffs across players may be controversial.

Once the evolutionary dynamics of the game is precisely defined, the emphasis in EGT is placed on studying which behavioral phenotypes (i.e. strategies) are stable under such evolutionary dynamics, and how such evolutionary stable states may be reached (Weibull 1995). Despite having its origin in biology, the basic ideas behind EGT –that successful strategies tend to spread more than unsuccessful ones, and that fitness is frequency-dependent– have extended well beyond the biological realm. In fact, nowadays there are a number of formal results that link several solution concepts in EGT (which were conceived as the result of the workings of evolution) with solution concepts in CGT (which were derived as the outcome of players’ introspective rational thinking) (see e.g. chapter 10 in Vega-Redondo 2003).

Learning game theory (LGT): As in classical game theory, players’ goal in most LGT models is to obtain the maximum possible payoff. However, LGT abandons the demanding assumptions of classical game theory on players’ rationality and beliefs, and assumes instead that players learn over time about the game and about the behavior of others (e.g. through reinforcement, imitation, or belief updating).

The process of learning in LGT can take many different forms, depending on the available information, the available feedback, and the way these are used to modify behavior. The assumptions made in these regards give rise to different models of learning. In most models of LGT, and in contrast with CGT, players use the history of the game to decide what action to take. In the simplest forms of learning (e.g. reinforcement or imitation) this link between acquired information and action is direct (e.g. in a stimulus-response fashion); in

more sophisticated learning, players use the history of the game to form expectations or beliefs about the other players' behavior, and they then react optimally to these inferred expectations.

The following is a brief list of some models of learning studied in LGT. We present these in ascending order of sophistication according to the amount of information that players use and their computational capabilities.

Reinforcement learning

Reinforcement learners rely on their experience to choose or avoid certain actions based on their immediate consequences. Actions that led to satisfactory outcomes in the past tend to be repeated in the future, whereas choices that led to unsatisfactory experiences are avoided. In general, reinforcement learners do not use more information than the immediately received payoff, which is used to adjust the probability of the conducted action accordingly. Reinforcement learners may well be presumed unaware of the strategic nature of the game.

Learning by imitation

Imitation occurs whenever a player –the imitator– adopts the strategy of some other player –the imitated. The definition of a particular imitation rule dictates when and how imitation takes place. Some models prescribe that players receive an imitation opportunity with some fixed independent probability; in other models the revision opportunity is triggered by some internal event (e.g. player's average payoff falling down below a certain threshold). When given the chance to revise her strategy, the imitator selects one other player to imitate; this selection is most often influenced by the payoff obtained by the other players in previous rounds, and it often leaves room for experimentation (i.e. adoption of a randomly selected strategy).

Interestingly, models of learning by imitation and evolutionary models are closely related: one can always understand an evolutionary model in learning terms, by re-interpreting the death-birth process as a strategy revision-imitation process conducted by immortal individuals. With this view in mind, one could argue that LGT actually encompasses EGT. However, if not essential in purely formal terms, the distinction between EGT and this particular subset of LGT is clear in the way models are formulated and interpreted, and also in the type of formal models studied in each discipline. A common difference between imitation models in the LGT literature and models in EGT is the level at which dynamic processes are defined. Models in LGT describe how players *individually* adapt through learning, and it is this learning process that is *explicitly* modeled. By contrast, many models in EGT are aggregate in the sense that they impose a dynamic process at the population level, abstracting from the micro-foundations that could give rise to such population dynamics.

Static perceptions and myopic response

In this family of learning models, each player is assumed to know the payoff that she would receive in each possible outcome of the game and the actions that every player selected in the immediate past. When making her next decision, every player assumes that every other player will keep her current action unchanged (i.e. static perception of the environment). Working under such assumption, each player can identify the set of strategies that would lead to an improvement in her current payoff. At this point different models posit different rules. *Better-response rules* assume that players select one of these payoff-improving strategies probabilistically, while the more demanding *best-response rule* assumes that players select a strategy which would have yielded the highest payoff. Thus, in these models players assume that their environment is static and

deterministic, and respond to it in a myopic fashion, i.e. ignoring the implications of current choices on future choices and payoffs.

Fictitious play

As in the previous class of models, players in fictitious play (FP) models are assumed to have a certain model of the situation and decide optimally on the basis of it. The higher level of sophistication introduced in FP models concerns the (still stationary) model of the environment that players hold. An FP player assumes that each of her counterparts is playing a certain mixed strategy, and her estimation of this mixed strategy is equal to the frequency with which the counterpart has selected each of her available actions up until that moment. Thus, instead of considering the actions taken by every other player only in the immediately preceding time-step (as in the models explained in the previous section), FP players implicitly take into account the whole history of the game. After forming her beliefs about every other player's strategy in such a frequentist manner, an FP player responds optimally (and myopically) to such beliefs.

Rational learning

The most sophisticated model of learning in LGT is often labeled “rational learning” – see Kalai and Lehrer (1993). Players in this model are assumed to be fully aware of the strategic context they are embedded in. They are also assumed to have a set of subjective beliefs over the behavioral strategies of the other players. Informally, the only assumption made about such beliefs is that players cannot be “*utterly surprised*” by the course of the play, i.e. players must assign a strictly positive probability to any strategy profile that is coherent with the history of the game. Finally, players are assumed to respond optimally to their beliefs with the objective of maximizing the flow of future payoffs discounted at a certain rate.

Note that all the models presented above embody some sort of rationality in the sense that players try to attain the maximum possible payoff, given certain constraints on the available information and on the formation of beliefs about other players. There are, however, other learning models in the literature which are not built upon the assumption that players respond optimally, but rather attempt to describe how humans actually play games (which does not always seem to be very rational). These learning models are inspired by empirical evidence and by research in cognitive science, and many of them are collected under the umbrella of Behavioral Game Theory (BGT). BGT “*is about what players actually do. It expands analytical theory by adding emotion, mistakes, limited foresight, doubts about how smart others are, and learning to analytical game theory*” (Camerer, 2003). Models in BGT are assessed according to how well they fit empirical (mostly experimental) data.

Important Scientific Research and Open Questions

Both EGT and LGT are very lively fields of research. Understandably, much research on mainstream EGT is founded on assumptions made to ensure that the resulting models are mathematically tractable. Thus, much research assumes infinite and homogeneous populations where players use one of a finite set of strategies and are randomly matched to play a 2-player symmetric game. The analysis of richer and more realistic systems (which consider e.g. finite populations, multi-player games, simultaneous mutations, and structured populations) has advanced a lot in recent years and is benefiting as well considerably from the advancement of computer simulation.

As for LGT, a common weakness of most models in the literature is that they almost invariably assume that every player in the game follows the same decision-making rule. This seems to be the natural first step in exploring the implications of a decision-making rule; however, it is clear that in many of these models the

observed dynamics are very dependent on the fact that the game is played among “cognitive clones”. Confronting the investigated learning algorithm with other decision-making rules seems to be a promising second step in LGT studies. As a matter of fact, the inclusion of different learning rules within the same model opens up a promising avenue of interaction between LGT and EGT: the evolution of learning rules, i.e. what type of learning rules may survive and spread in an evolutionary context?

There is also a lot of research to be done on the relation between CGT, EGT and LGT. A topic of much interest lies in studying the conditions under which the solution concepts derived in each of these fields coincide, e.g. When does a certain learning rule converge to a Nash equilibrium?, Under what conditions evolution favors rational behavior? Are the dynamics of a certain evolutionary process formally equivalent to those obtained when the game is played by individuals who learn to play the game?

Another lively area of research concerns the *computational complexity* of the problems we encounter in game theory. Any problem posed in the context of game theory requires some sort of computation. This does not only apply to LGT (where it is clear that learning algorithms have to compute the strategy to use as a function of the available information and feedback) but to the whole field of game theory in general. The identification of best-response strategies, Nash equilibria, evolutionary stable strategies, or of any other solution concept, are computational problems in the sense that they require an algorithm, a procedure to compute the results. Computational complexity theory studies the inherent difficulty of computational problems, i.e. how hard the problems are, measured as a function of the amount of computational resources needed to solve them. Knowing the computational complexity of a problem can be of utmost significance, and it can crucially influence the applicability of its solution.

Finally, there is clearly a lot to gain from the interaction of game theory with other disciplines. Traditionally, game theory has developed almost entirely from introspection and theoretical concerns. Whilst the work developed in game theory up until now has proven to be tremendously useful, it seems clear that game theory will not fulfill all its potential as a tool to analyze real-world social interactions unless greater attention is paid to empirical evidence and concrete real-world problems. Empirical research (both experimental and field work) can also suggest exciting and relevant avenues where theoretical research may be most needed. In this way, empirical and theoretical work can usefully drive, shape, and benefit from each other.

References

- Camerer, C. (2003). Behavioral Game Theory: Experiments on Strategic Interaction. New York: Russell Sage Foundation.
- Hamilton, W. D. (1967). Extraordinary sex ratios. *Science* 156(3774), 477-488.
- Kalai, E. & Lehrer, E. (1993). Rational Learning Leads to Nash Equilibrium. *Econometrica* 61(5), 1019-1045.
- Vega-Redondo, F. (2003). Economics and the Theory of Games. Cambridge University Press.
- Von Neumann, J. & Morgenstern, O. (1944). Theory of Games and Economic Behavior. Princeton University Press.
- Weibull, J. W. (1995). Evolutionary Game Theory. Cambridge, MA: MIT Press.